

Solving systems of linear equations on a quantum computer

Stefanie Barz¹, Ivan Kassal^{2,3}, Martin Ringbauer^{1,*}, Yannick Ole Lipp¹, Borivoje Dakić¹,
Alán Aspuru-Guzik², Philip Walther¹

¹ Faculty of Physics, University of Vienna, Boltzmanngasse 5, 1090 Vienna, Austria

² Department of Chemistry and Chemical Biology, Harvard University, Cambridge MA 02138, United States

³ Centre for Engineered Quantum Systems, Centre for Quantum Computing and Communication Technology, and School of Mathematics and Physics, The University of Queensland, St Lucia QLD 4072, Australia

* Present address: Centre for Engineered Quantum Systems,
Centre for Quantum Computing and Communication Technology,
and School of Mathematics and Physics, The University of Queensland, St Lucia QLD 4072, Australia

Systems of linear equations are used to model a wide array of problems in all fields of science and engineering. Recently, it has been shown that quantum computers could solve linear systems exponentially faster than classical computers [1], making for one of the most promising applications of quantum computation [2]. Here, we demonstrate this quantum algorithm by implementing various instances on a photonic quantum computing architecture. Our implementation involves the application of two consecutive entangling gates on the same pair of polarisation-encoded qubits. We realize two separate controlled-NOT gates where the successful operation of the first gate is heralded by a measurement of two ancillary photons. Our work thus demonstrates the implementation of a quantum algorithm with high practical significance as well as an important technological advance which brings us closer to a comprehensive control of photonic quantum information [3].

Systems of linear equations play an important role in various fields, ranging from natural science and engineering to medicine and social science. The ability to solve such systems underpins many modern technologies, including traffic flow analysis, computer tomography, and weather forecasting. Although systems of linear equations are an old problem—the two-dimensional version was investigated by the ancient Babylonians—it was only in 1811 that Gauss developed a general algorithm for solving them, Gaussian elimination [4]. Today, as the sizes of data sets are growing, Gaussian elimination is too slow and more advanced methods are needed. An increasing demand for a detailed understanding of larger and larger systems pushes the tools of classical computation to their limits. Modern data sets can be so enormous that finding a solution to a system of linear equations can be prohibitive even for the latest supercomputers.

Quantum computers have attracted tremendous interest because they could outperform classical computers at certain tasks [5]. For a classical computer, the number of computational steps needed to solve a linear system is at least proportional to the number of variables. By contrast, the new quantum algorithm could, in some cases, make the computational time proportional to only the logarithm of the number of variables [1]. An important difference is that the quantum algorithm calculates the expectation value of an operator associated with the solution rather than the solution itself. Here, we demonstrate this algorithm using two full controlled-NOT (CNOT) gates acting on two qubits to determine the solution of a two-dimensional system of linear equations. We show various instances of the algorithm for systems of linear equations with different characteristics.

THEORY

Solving a linear system of equations, given a matrix A and a vector b , means finding the vector x such that $Ax = b$. If we rescale the vectors to $\|b\| = \|x\| = 1$, we can represent them as quantum states $|b\rangle$ and $|x\rangle$, and our initial task becomes finding the state $|x\rangle$ such that

$$A|x\rangle = |b\rangle. \quad (1)$$

The solution we seek is

$$|x\rangle = \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|}. \quad (2)$$

We assume, without loss of generality [1], that A is an $N \times N$ Hermitian matrix with eigenbasis $\{|u_j\rangle\}$ and eigenvalues $\{\lambda_j\}$, and is rescaled so that $0 < \lambda_j < 1$. The state $|b\rangle$ can be expanded in the eigenbasis, $|b\rangle = \sum_{j=1}^N \beta_j |u_j\rangle$, and we aim to prepare, up to normalization,

$$|x\rangle = \sum_{j=1}^N \beta_j \frac{1}{\lambda_j} |u_j\rangle. \quad (3)$$

The algorithm consists of three main steps. Here, we describe the basic idea of the algorithm; a more detailed discussion can be found in [1] and the Appendix).

The first step is to apply phase estimation, a general procedure for decomposing quantum states in a particular basis [6–8]. For this, we add an additional “eigenvalue register” of m qubits to our system, each initialized in the state $|0\rangle$. Phase estimation then transforms $|b\rangle|0\rangle^{\otimes m}$ into $\sum_{j=1}^N \beta_j |u_j\rangle |\lambda_j\rangle$,

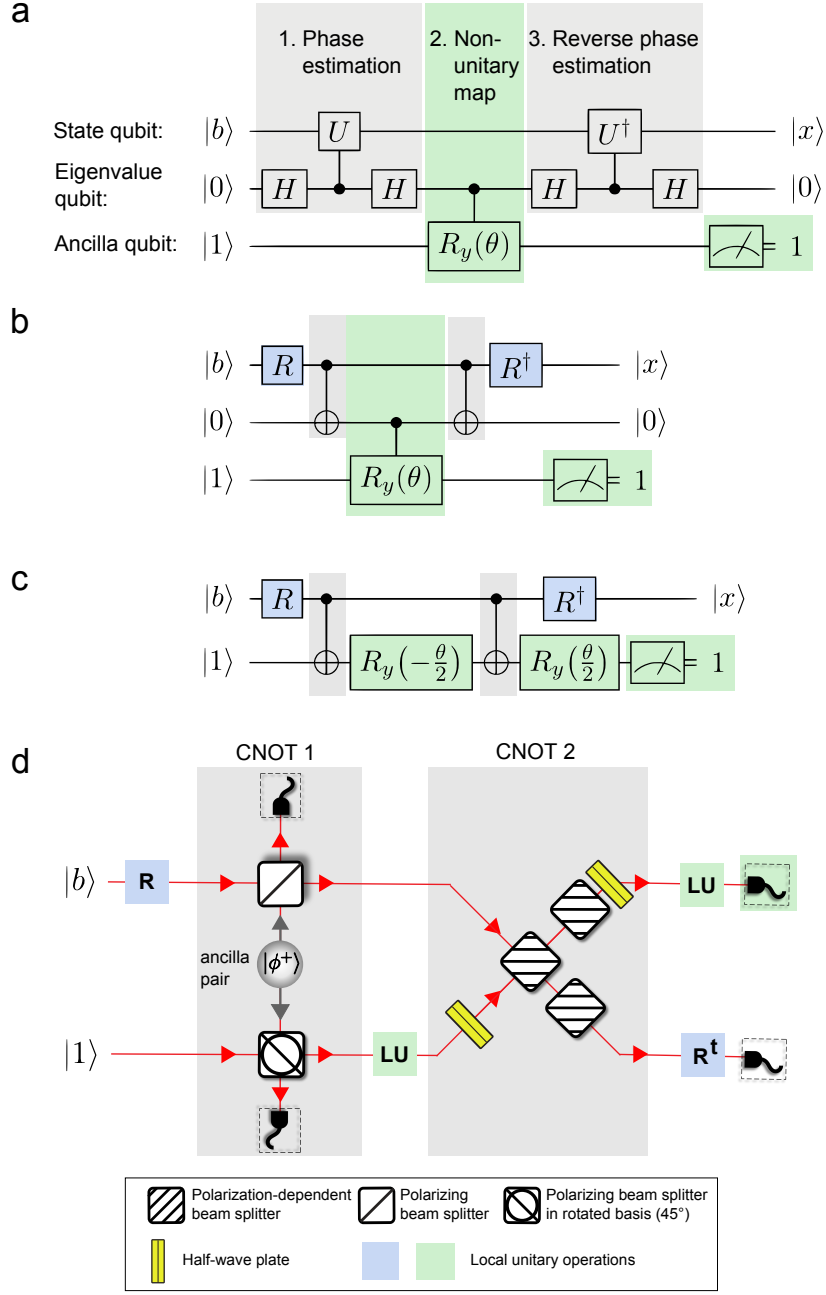


FIG. 1: The simplest case of the quantum algorithm for solving systems of linear equations. Given a Hermitian matrix A and input $|b\rangle$, outputs $|x\rangle = A^{-1}|b\rangle / \|A^{-1}|b\rangle\|$ if the ancilla qubit is measured to be 1. **(a)** The complete circuit, as derived in the text, with $U = \exp(2\pi i 2^{n-1} A)$ and $\theta = -2 \arccos(\lambda_1/\lambda_2)$, where $\lambda_{1,2}$ are the eigenvalues of A and the integer n depends on the eigenvalues (see Appendix). **(b)** The local unitary R diagonalises A , $A = R^\dagger \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} R$. For the algorithm to work perfectly with this many qubits, $\lambda_{1,2}$ must be such that $U = R^\dagger \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} R = R^\dagger Z R$. This circuit reflects this simplification. **(c)** Optimised circuit. The middle qubit can be completely removed, and the controlled-rotation decomposed, to give the final circuit (see Appendix for details). $|b\rangle$ and R are arbitrary. **(d)** Experimental implementation of the circuit shown in c. The local unitary operations are implemented with the help of a combination of half-wave plates and quarter-wave plates.

where the eigenvalues $|\lambda_j\rangle$ are stored in the eigenvalue register to a precision of m binary digits.

The second step is to implement the nonunitary map $|\lambda_j\rangle \rightarrow \lambda_j^{-1}|\lambda_j\rangle$. For this, we introduce an additional

“ancilla qubit” initially in the state $|1\rangle$. Depending on the value of $|\lambda_j\rangle$ in the eigenvalue register, we implement a controlled- $R_y(\theta_j)$ rotation on the ancilla qubit. Here, $R_y(\theta) = \exp(-i\theta\sigma_y/2)$ and σ_y denotes the usual Pauli matrix. With




FIG. 2: Experimental setup. Shown is the experimental implementation of two concatenated CNOT gates. The input is set by a polariser, which can be followed by a local unitary operation (LU). The two gates are connected by fibers. Different matrices A can be implemented by adapting the LUs, and different states $|b\rangle$ by adapting the input state. The figure shows the most general case of two concatenated CNOT gates, combined with general LUs. For the implementation of the algorithm, we chose some LUs to be the identity and obtained the case shown in Figure 1d.

$\theta_j = -2 \arccos(C/\lambda_j)$, the state of our system becomes

$$\sum_{j=1}^N \beta_j |u_j\rangle |\lambda_j\rangle \left(\sqrt{1 - \frac{C^2}{\lambda_j^2}} |0\rangle + \frac{C}{\lambda_j} |1\rangle \right), \quad (4)$$

where $C \leq \min_j |\lambda_j|$.

Finally, the third step is to run phase estimation in reverse to uncompute $|\lambda_j\rangle$, giving

$$\sum_{j=1}^N \beta_j |u_j\rangle \left(\sqrt{1 - \frac{C^2}{\lambda_j^2}} |0\rangle + \frac{C}{\lambda_j} |1\rangle \right). \quad (5)$$

We measure the ancilla qubit, and if we observe a 1 we will have prepared $|x\rangle$ in the state register. If we know the eigenvalues, we can maximise the success probability by choosing the largest possible C , $C = \min_j |\lambda_j|$.

The runtime of the algorithm is $\tilde{O}(\log(N)s^2\kappa^2/\varepsilon)$, where s is the sparsity of the matrix, κ its condition number, and ε the acceptable error [1]. Furthermore, $\tilde{O}(\kappa^2)$ can be reduced to $\tilde{O}(\kappa)$ with amplitude amplification [9]. The best classical algorithms require $O(Ns\sqrt{\kappa}\log(1/\varepsilon))$ time, meaning that at constant s , κ , and ε , the quantum algorithm is exponentially faster. On a quantum computer, determining all amplitudes in a quantum state scales exponentially with system size; instead, the strength of the algorithm lies in the determination of expectation values $\langle x|\hat{M}|x\rangle$ of some operator \hat{M} .

The simplest case involves three qubits: one state qubit, one eigenvalue qubit, and the ancilla qubit. The state and eigenvalue qubits replace larger registers which would be needed in a general implementation of the algorithm. Whereas using one state qubit means that $|b\rangle$ is a two-vector and A a 2×2 matrix, using one eigenvalue qubit imposes more subtle restrictions. With one eigenvalue qubit, only a single binary digit of the eigenvalues is computed by the phase estimation, meaning that for the algorithm to work perfectly, it must be possible to distinguish the two eigenvalues with a single digit. Consequently, we choose the two eigenvalues to be of the form $0.\bar{a}0$ and $0.\bar{a}1$, where \bar{a} is a sequence of binary digits.

The complete circuit for the algorithm as described is given in Figure 1, which also outlines the procedure for optimising the circuit to require only two qubits and two consecutive CNOTs acting on them (Figure 1c). The circuit depends on the eigenvalues of A , the unitary R that diagonalizes it,

$$A = R^\dagger \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} R, \quad (6)$$

and the input state $|b\rangle$. The algorithm succeeds—i.e. the ancilla qubit is measured in the state $|1\rangle$ —with probability $(\lambda_1/\lambda_2)^2$.

In our implementation, we choose different sets of eigenvalues $\Lambda = \{\lambda_1, \lambda_2\}$. The simplest case is $\bar{a} = \bar{1}$, giving

the eigenvalues $\lambda_1 = 0.10 = \frac{1}{2}$ and $\lambda_2 = 0.11 = \frac{3}{4}$. In this case phase estimation must read out the second bit of the eigenvalues. The procedure is analogous for the other sets of eigenvalues that we implemented (see Appendix).

EXPERIMENTAL REALIZATION

We implemented the algorithm using polarisation-encoded photonic qubits (Figures 1d and 2), where $|0\rangle$ and $|1\rangle$ denote horizontal and vertical polarisation, respectively [10, 11]. Our implementation uses two different types of photonic CNOT gates [12–18].

The first CNOT gate realizes the gate operation in a heralded manner [19, 20]. It requires an entangled ancilla photon-pair and a measurement of two ancilla modes to herald that the gate has worked correctly. If two photons are registered in the ancilla modes, the gate has been successful without the need for a verification of the output state. Since the output photons need not to be measured, the application of a second CNOT gate is possible.

Our second CNOT is implemented in a destructive way, where a coincident measurement of the output photons signals the correct gate operation. The basic element of this destructive CNOT gate is a polarisation-dependent beam splitter (PDBS) which has a different transmission coefficient T for horizontally polarised light ($T_H = 1$) as for vertically polarised light ($T_V = 1/3$) [14]. If two vertically-polarised photons are reflected at this PDBS, they acquire a phase shift of π . Two successive PDBSs with the opposite splitting ratios then equalize the output amplitudes. This setup, in combination with two half-wave plates (HWPs) (see Figure 1d) implements a destructive CNOT gate. The gate has been successful if one photon is measured in each output mode.

Combining these photonic CNOT gates with local unitary operations allows us to implement the circuit shown in Figure 1. In our setup, we implement these local unitary operations with the help of a combination of quarter-wave plates (QWPs) and HWPs (see Figures 1d and 2). A detailed description of our experimental setup can be found in the Appendix.

IMPLEMENTATION OF THE ALGORITHM

We have implemented various instances of the algorithm, where we varied both the matrix A and the state $|b\rangle$. The matrix can be modified by tuning the local operation R (in which case the eigenvalues stay the same) or by adapting θ , which alters the eigenvalues of the matrix. The detection of the ancilla qubit in the state $|1\rangle$ announces a successful run of the algorithm and the preparation of the output qubit in the state $|x\rangle$.

Figure 3 shows the results of a sample run of the algorithm. Two different matrices A were implemented by using different local operations R_1 and R_2 . The eigenvalues were $\Lambda = \{\frac{1}{2}, \frac{3}{4}\}$ in both cases. We then choose state vectors $|b_1\rangle$

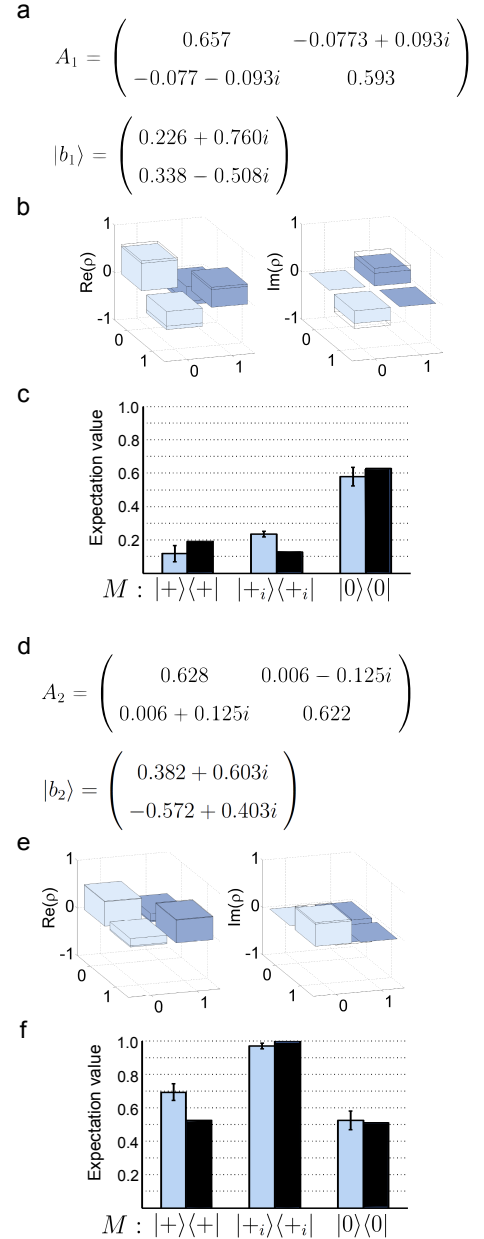


FIG. 3: Experimental results. (a), (d) The figure shows two different systems of linear equations, depicted by the matrices A_1 and A_2 as well as the state vectors $|b_1\rangle$ and $|b_2\rangle$. (b), (e) The reconstructed density matrices of the experimentally obtained output state $|x\rangle$ are shown. These density matrices are obtained by choosing the local operations $R_1 = R_x(\frac{11}{15}\pi).R_y(\frac{3}{8}\pi)$ (b), and $R_2 = R_x(\frac{89}{60}\pi).R_y(-\frac{3}{8}\pi)$ (e). For both matrices, we choose the eigenvalues to be $\lambda_1 = \frac{1}{2}$ and $\lambda_2 = \frac{3}{4}$ by implementing $R_y(\theta)$ as described in the main text. The fidelities of the reconstructed density matrices are 0.953 ± 0.026 (b) and 0.976 ± 0.010 (e). The wireframe shows the theoretical prediction. (c), (f) The quantum algorithm is based on determining the expectation value $\langle x|\hat{M}|x\rangle$ of some operator \hat{M} with respect to the output state $|x\rangle$. Therefore, we also show the experimentally determined (blue) and theoretical (black) expectation values of several operators \hat{M} . We choose the operator \hat{M} to be the projection on the states $|0\rangle$, $|+\rangle$, and $|+_i\rangle$, respectively, with $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, and $|+_i\rangle = (|0\rangle + i|1\rangle)/\sqrt{2}$.

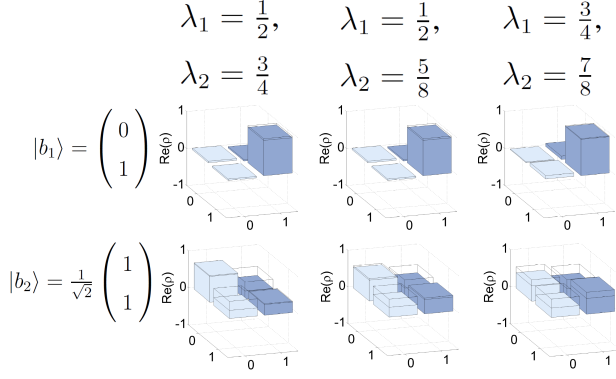


FIG. 4: The figure shows the solution of the system of linear equation for matrices with different eigenvalues. Experimentally, these are obtained by implementing different values of θ (see Appendix for details). For all matrices we run the algorithm for two input states $|b_1\rangle = |1\rangle = (0, 1)$ and $|b_2\rangle = |+\rangle = (1, 1)/\sqrt{2}$. We achieve fidelities of 0.957 ± 0.010 , 0.961 ± 0.013 , 0.981 ± 0.009 for the input state $|b_1\rangle$ (upper row from left to right) and 0.778 ± 0.031 , 0.773 ± 0.027 , 0.832 ± 0.031 for the input state $|b_2\rangle$ (lower row, left to right).

and $|b_2\rangle$ by setting different input states in our experiment (see Figure 3a). We run the algorithm and analyse the output state $|x\rangle$ via quantum state tomography. The density matrices of the resulting output states $|x\rangle$ are shown in Figure 3c. As the algorithm relies on the determination of the expectation values $\langle x|\hat{M}|x\rangle$, the figure also shows expectation values with several operators \hat{M} (see Figure 3d).

Furthermore, we implemented the algorithm for a series of matrices A with three different sets of eigenvalues $\Lambda_1 = \{\frac{1}{2}, \frac{3}{4}\}$, $\Lambda_2 = \{\frac{1}{2}, \frac{5}{8}\}$, and $\Lambda_3 = \{\frac{3}{4}, \frac{7}{8}\}$ (see Figure 4). As explained in detail in the Appendix, the performance of the algorithm depends on the state $R|b\rangle$ that enters the gate. In order to analyse this behaviour, we chose two different input states, $|b_1\rangle = |1\rangle$ and $|b_2\rangle = |+\rangle$ for each set of eigenvalues, while keeping R equal to the identity matrix. The six resulting density matrices are shown in Figure 4. We achieve fidelities of up to 0.981 ± 0.009 for $|b_1\rangle$, and 0.832 ± 0.031 for $|b_2\rangle$. This difference in fidelities arises due to the influence of higher-order emissions, which depends on the state $R|b\rangle$, as discussed in the Appendix. Additional data for a set of different input states and the three different choices of eigenvalues is shown in the Appendix.

DISCUSSION

The results presented here include the implementation of the simplest case of the quantum algorithm for solving systems of linear equations, as well as the concatenation of two entangling gates acting on the same photonic qubits. We anticipate that increasing technological capabilities, including the implementation of more than two consecutive CNOTs, will allow the extension of the algorithm both to larger sys-

tems and to more precise and full phase estimation. In particular, the ability to do arithmetic operations with quantum gates will enable the calculation of inverses (and thus the map $|\lambda_j\rangle \rightarrow \lambda_j^{-1} |\lambda_j\rangle$) on the fly. Even though current quantum computations in optical systems are proof-of-principle demonstrations, important insights can be obtained for future realisations in larger systems [21–23]. We are likewise hopeful that our demonstration of this algorithm will enable future implementations of other, equally important algorithms that use it as a subroutine, including quantum algorithms for solving nonlinear differential equations [24] and quantum data fitting [25].

-
- [1] A. Harrow, A. Hassidim, and S. Lloyd, Phys. Rev. Lett. **103**, 150502 (2009).
 - [2] A. Childs, Nature Physics **5**, 861 (2009).
 - [3] E. Knill, R. Laflamme, and G. J. Milburn, Nature **409**, 46 (2001).
 - [4] I. Kleiner, *A history of abstract algebra* (Birkhäuser, 2007).
 - [5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
 - [6] A. Kitaev, Russ. Math. Surv. **52**, 1191 (1997).
 - [7] D. Berry, G. Ahokas, R. Cleve, and B. Sanders, Commun. Math. Phys. **270**, 359 (2007).
 - [8] A. Childs, Commun. Math. Phys. **294**, 581 (2010).
 - [9] A. Ambainis, arXiv:1010.4458 (2010).
 - [10] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, Rev. Mod. Phys. **79**, 135 (2007).
 - [11] J. O’Brien and J. Akira Furusawa, Vučković, Nature Photon. **3**, 687 (2009).
 - [12] J. L. O’Brien, G. J. Pryde, A. G. White, T. C. Ralph, and D. Branning, Nature **426**, 264 (2003).
 - [13] N. K. Langford, T. J. Weinhold, R. Prevedel, K. J. Resch, A. Gilchrist, J. L. O’Brien, G. J. Pryde, and A. G. White, Phys. Rev. Lett. **95**, 210504 (2005).
 - [14] N. Kiesel, C. Schmid, U. Weber, R. Ursin, and H. Weinfurter, Phys. Rev. Lett. **95**, 210505 (2005).
 - [15] R. Okamoto, H. F. Hofmann, S. Takeuchi, and K. Sasaki, Phys. Rev. Lett. **95**, 210506 (2005).
 - [16] X. Zhou, T. Ralph, P. Kalasuwan, M. Zhang, A. Peruzzo, B. Lanyon, and J. O’Brien, Nature Commun. **2**, 413 (2011).
 - [17] A. Crespi, R. Ramponi, R. Osellame, L. Sansoni, I. Bongioanni, F. Sciarrino, G. Vallone, and P. Mataloni, Nature Commun. **2**, 566 (2011).
 - [18] L. Sansoni, F. Sciarrino, G. Vallone, P. Mataloni, A. Crespi, R. Ramponi, and R. Osellame, Phys. Rev. Lett. **105**, 200503 (2010).
 - [19] T. Pittman, B. Jacobs, and J. Franson, Phys. Rev. A **64**, 062311 (2001).
 - [20] S. Gasparoni, J.-W. Pan, P. Walther, T. Rudolph, and A. Zeilinger, Phys. Rev. Lett. **93**, 020504 (2004).
 - [21] B. Lanyon, J. Whitfield, G. Gillett, M. Goggin, M. Almeida, I. Kassal, J. Biamonte, M. Mohseni, B. Powell, M. Barbieri, et al., Nature Chem. **2**, 106 (2010).
 - [22] X. Zhou, P. Kalasuwan, T. Ralph, and J. O’Brien, arXiv:1110.4276 (2011).
 - [23] E. Martín-López, A. Laing, T. Lawson, R. Alvarez, X. Zhou, and J. O’Brien, Nature Photon. (2012).
 - [24] S. K. Leyton and T. J. Osborne, arXiv:0812.4423 (2008).

- [25] N. Wiebe, D. Braun, and S. Lloyd, Phys. Rev. Lett. **109**, 050505 (2012).
 [26] B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O'Brien, A. Gilchrist, and A. G. White, Nature Phys. **5**, 134 (2009).
 [27] P. G. Kwiat, E. Waks, A. G. White, I. Appelbaum, and P. H. Eberhard, Phys. Rev. A **60**, R773 (1999).

Acknowledgments We thank Frank Verstraete for valuable discussions and Thomas Lindner for assistance in the laboratory. I.K. was supported by a UQ Postdoctoral Research Fellowship and acknowledges support from the Australian Research Council Centres of Excellence for Engineered Quantum Systems (project CE110001013) and Quantum Computation and Communication Technology (project CE110001027). A.A.G. acknowledges support from the Hughes Research Laboratory (grant M1144-201167-DS) and the Air Force Research Office (contract FA9550-12-1-0046 and 10323836-SUB) and thanks the Corning, Sloan and Dreyfus foundations for their support. P.W. acknowledges support from the European Commission, Q-ESSENCE (No. 248095), QUILMI (No. 295293) and the ERA-Net CHISTERA project QUASAR, the John Templeton Foundation, the Vienna Center for Quantum Science and Technology (VCQ), the Austrian Nano-initiative NAP Platon, the Austrian Science Fund (FWF) through the SFB FoQuS (No. F4006-N16), START (No. Y585- N20) and the doctoral programme CoQuS, the Vienna Science and Technology Fund (WWTF) under grant ICT12-041, and the Air Force Office of Scientific Research, Air Force Material Command, United States Air Force, under grant number FA8655-11-1-3004.

APPENDIX

Theory

Our work differs from the original proposal in [1] in several modifications that are needed to implement the algorithm with a limited number of qubits. Some of these modifications were mentioned in the main text, and we elaborate on them here, essentially describing the simplifications that are used to transform the circuit in Figure 1a to the circuit in Figure 1c.

As described in the text, the most general circuit involves phase estimation, a controlled R_y rotation, and a reverse phase estimation (for a description of phase estimation in general, see Sec. 5.2 of [5]). If we restrict ourselves to one state qubit and one eigenvalue qubit, phase estimation is simply a controlled unitary between two Hadamard gates on the eigenvalue qubit (see Figure 1a).

The unitary U that is chosen depends on what binary digit of the eigenvalue needs to be read out. Because we have assumed that A is scaled in such a way that its eigenvalues lie in

the range $(0, 1)$, we can express them as follows:

$$\lambda_1 = 0.\overline{x_{1,1}x_{1,2}x_{1,3}x_{1,4}} \dots \quad (7)$$

$$\lambda_2 = 0.\overline{x_{2,1}x_{2,2}x_{2,3}x_{2,4}} \dots, \quad (8)$$

where the $\bar{x}_{i,j}$ are binary digits, 0 or 1. If we wish to read out the n^{th} digit, we must choose $U = \exp(2\pi i 2^{n-1} A)$. Note its action on the eigenvectors:

$$U|u_j\rangle = \exp(2\pi i 2^{n-1} \lambda_j) |u_j\rangle \quad (9)$$

$$= \exp(2\pi i 2^{n-1} 0.\overline{x_{j,1}x_{j,2}x_{j,3}x_{j,4}} \dots) |u_j\rangle \quad (10)$$

$$= \exp(2\pi i 0.\overline{x_{j,n}x_{j,n+1}x_{j,n+2}} \dots) |u_j\rangle \quad (11)$$

$$\approx \exp(2\pi i 0.\overline{x_{j,n}}) |u_j\rangle \quad (12)$$

$$= (-1)^{\overline{x_{j,n}}} |u_j\rangle. \quad (13)$$

The approximation of neglecting all digits past \bar{x}_n introduces errors to the procedure, which would be mitigated by using additional eigenvalue qubits. Without the additional qubits, we avoid the error by assuming, as in the main text, that the eigenvalues are of the form $0.\bar{a}0$ and $0.\bar{a}1$, differing only at the n^{th} digit.

The state and eigenvalue qubits, initialized to $|b\rangle_S|0\rangle_E$, transform thus:

$$|b\rangle_S|0\rangle_E = \sum_{j=1}^2 \beta_j |u_j\rangle |0\rangle \xrightarrow{H} \sum_{j=1}^2 \beta_j |u_j\rangle \frac{|0\rangle + |1\rangle}{\sqrt{2}} \quad (14)$$

$$\xrightarrow{\text{cont.}-U} \sum_{j=1}^2 \beta_j |u_j\rangle \frac{|0\rangle + (-1)^{\overline{x_{j,n}}} |1\rangle}{\sqrt{2}} \quad (15)$$

$$\xrightarrow{H} \sum_{j=1}^2 \beta_j |u_j\rangle |\bar{x}_{j,n}\rangle. \quad (16)$$

The two qubits are now entangled, with each eigenstate in the state register accompanied by the n^{th} digit of its eigenvalue.

As described in the main text, $\theta = -2 \arccos(\lambda_1/\lambda_2)$ in all cases. For example, if, as in Figure 3, $\lambda_1 = 0.10 = \frac{1}{2}$ and $\lambda_2 = 0.11 = \frac{3}{4}$, in which case we read out the second digit, we get $\theta = -1.682$. In the second column of Figure 4, $\lambda_1 = 0.100 = \frac{1}{2}$ and $\lambda_2 = 0.101 = \frac{5}{8}$, and so $\theta = -1.287$ (with the third digit read out). In the third column of Figure 4, $\lambda_1 = 0.110 = \frac{1}{2}$ and $\lambda_2 = 0.111 = \frac{5}{8}$, and so $\theta = -1.082$.

Once we assume that the eigenvalues differ in the n^{th} digit, the action of U is seen to be very simple: it is merely a Z gate in the eigenbasis of A . That is, it leaves the first eigenstate unchanged and adds a phase of -1 to the second. It can therefore be written as $U = R^\dagger Z R$, where R diagonalizes A . It is with this fact, and remembering that a controlled- Z conjugated with Hadamard gates gives a controlled-NOT, that one obtains the circuit in Figure 1b.

To get from Figure 1b to Figure 1c, we first observe that the eigenvalue qubit can be removed. After the gate R , the first CNOT transfers the state of the state register to the eigenvalue register. This then operates the controlled- $R_y(\theta)$ rotation. Instead of transferring the state first, it is equivalent to simply





FIG. 5: The figure shows the different states $R|b\rangle$ which we have chosen as input states to our circuit.

control the $R_y(\theta)$ gate directly from the state qubit, giving this circuit:



Controlled single-qubit rotations have already been implemented in linear optics [21, 22, 26], but here we follow a different approach. We decompose the controlled- R_y rotation using the general method described in Sec. 4.3 of [5], which immediately gives the circuit in Figure 1c.

Experimental setup

In our experiment, entangled photon pairs are produced by exploiting the emissions of a non-collinear type-II SPDC process [27]. For this, a mode-locked Mira HP Ti:Sa oscillator is pumped by a Coherent Inc. Verdi V-10 laser. The pulsed-laser output ($\tau = 200$ fs, $\lambda = 789$ nm, 76 MHz) is frequency-doubled using a 2 mm-thick lithium triborate (LBO) crystal, resulting in UV pulses of 0.75 W cw average. We achieve a stable source of UV pulses by translating the LBO to avoid optical damage to the anti-reflection coating of the crystal. The UV laser beam passes through a 2 mm-thick β -barium borate crystal, gets reflected, and passes through the crystal a second time (see Figure 2). The photons created during the first pass of the laser beam enter the first CNOT gate as the input (control and target) qubits. The state of these input qubits is modified using polarisers and additional local unitary gates that in principle allow for the creation of arbitrary input states. In our experiment, the control and target qubits are prepared in the states $R|b\rangle$ and $|1\rangle$, respectively. Thus, we absorb the local operation R in the preparation of the input state. The photons created when the laser passes through the crystal the second time act as the entangled ancilla photon pairs which is required for the first CNOT gate. For this, we align our setup such that the entangled state $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ is emitted [20].

The photons interfere at the polarising beam splitters (PBS) as shown in Figure 4. The PBS on the control side (target side)

is aligned such that it acts in the basis $\{|0\rangle, |1\rangle\}$ ($\{|+\rangle, |-\rangle\}$). The photons are the filtered spatially and spectrally with the help of narrow-band filters ($\Delta\lambda = 3$ nm) and by coupling them into single-mode fibers. A coincidence detection of the ancilla qubits in detectors 3 and 4 in the state $|-\rangle_3|1\rangle_4$ signals a successful gate operation.

The output photons in the modes I and II (see Figure 4) are then guided to the second CNOT gate. Wave plates before and after the second CNOT implement the local rotations $R_y(\theta/2)$, $R_y(\theta/2)$, R , and R^\dagger .

Photons are coupled out, pass the polarisation-dependent beam splitters (PDBSs), and are coupled to single-mode fibers again. The success of the second CNOT operation is determined by postselection on a coincidence detection in outputs 1 and 2. The algorithm succeeds if the target photon is detected in state $|1\rangle$. The output control qubit is in the state $|x\rangle$, which is analysed by using HWPs, QWPs and polarising beam splitters; and a full state tomography of the output state $|x\rangle$ is performed. Errors are obtained from a Monte Carlo routine assuming Poissonian counting statistics. These indicate a lower bound for the actual error that takes all the experimental imperfections into account. In our experiment, typical visibilities of the emitted Bell pairs are about 0.9 and higher-order emissions degrade the quality of our gate operations. The back-reflecting mirror is continuously moved back-and-forth to avoid any phase correlations between of the signal and the noise originating from higher-order photon emissions. Additionally, imperfect visibilities on the order of 0.85 to 0.9 of the quantum interference at the PBSs in the first gate and the PDBS in the second gate contribute to errors.

Experiment

As mentioned in the main text, the fidelity of the output state $|x\rangle$ depends on the state $R|b\rangle$ that effectively enters the control input of the first CNOT gate. The Figures 6, 7, and 8 show the characterization of the state $|x\rangle$ for various input states. Figure 5 depicts the different input states we have chosen for our analysis.

Our analysis shows that the fidelities of the obtained output

states vary from $(64.7 \pm 4.2)\%$ to $(98.1 \pm 0.9)\%$. These variations in fidelity arise due to the influence of higher-order emissions from spontaneous parametric down-conversion. These higher-order emissions can either occur when the beam passes the crystal the first time (“double-forward emission”) or when the beam passes the crystal the second time (“double-backward emission”).

If the input to the first CNOT gate is chosen to be $R|b\rangle = |0\rangle$ or $R|b\rangle = |1\rangle$, a double-forward emission can never lead to a fourfold coincidence and thus signal a wrong “successful” operation of the first gate. In these cases, the photons entering the control input to the gate will either both be reflected or both be transmitted at the first PBS. For all other inputs to the

gate, these higher-order emissions degrade the fidelity of the output state as demonstrated in our analysis.

The “double-backward emission” can in principle never lead to a fourfold coincidence because the photons can never split up due to quantum interference. However, due to the visibility of the entangled Bell pairs of 0.9, this quantum interference does not work perfectly. About 10% of our total counts arise from these events, which also influences the fidelity of the output state $|x\rangle$. However, for the input $R|b\rangle = |1\rangle$ the fidelity does not seem to be affected. In this case, the noise cannot be distinguished from the signal. The influence of the double-backward emission increases for the other input states and reaches a maximum for $R|b\rangle = |0\rangle$.

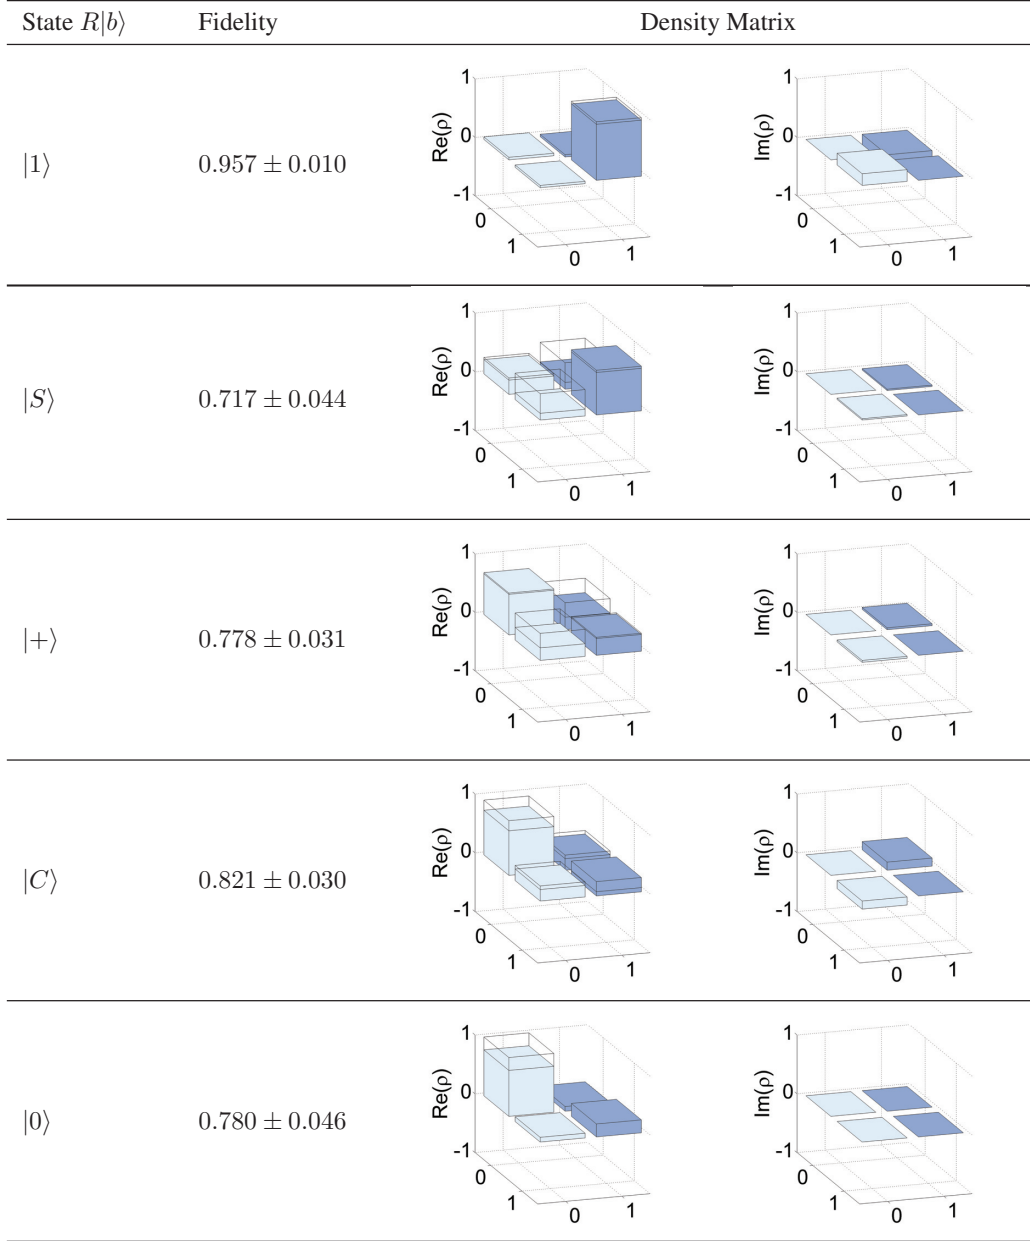


FIG. 6: The figure shows the different states $R|b\rangle$ which we have chosen as input states to our circuit.

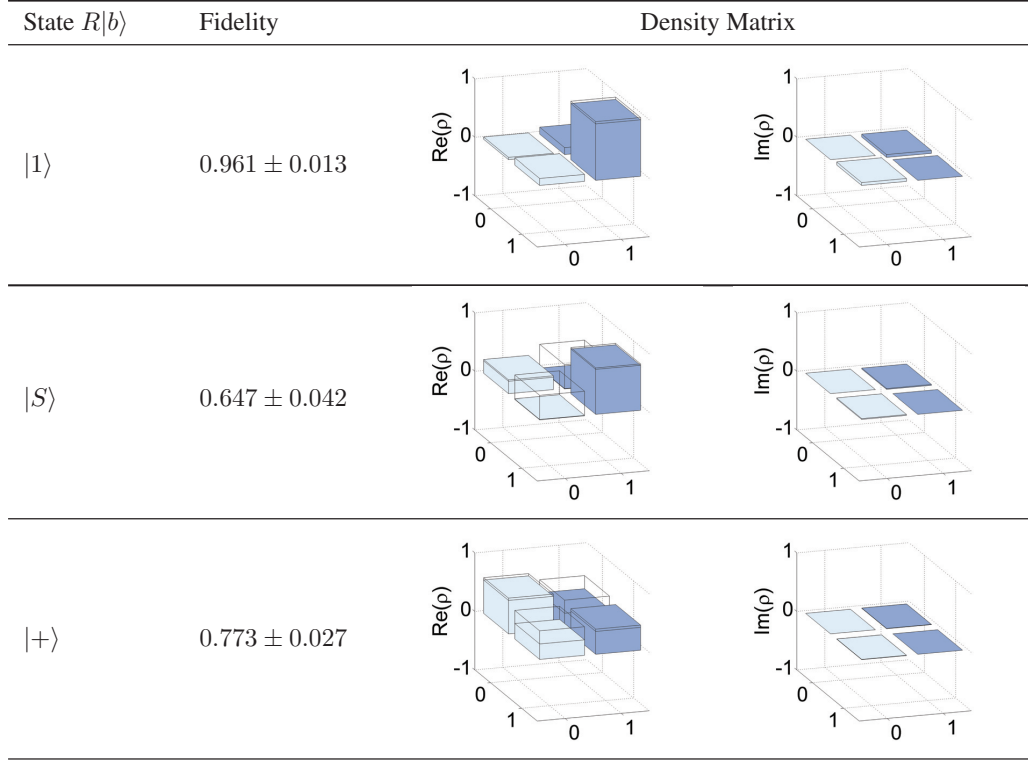


FIG. 7: The figure shows the different states $R|b\rangle$ which we have chosen as input states to our circuit.

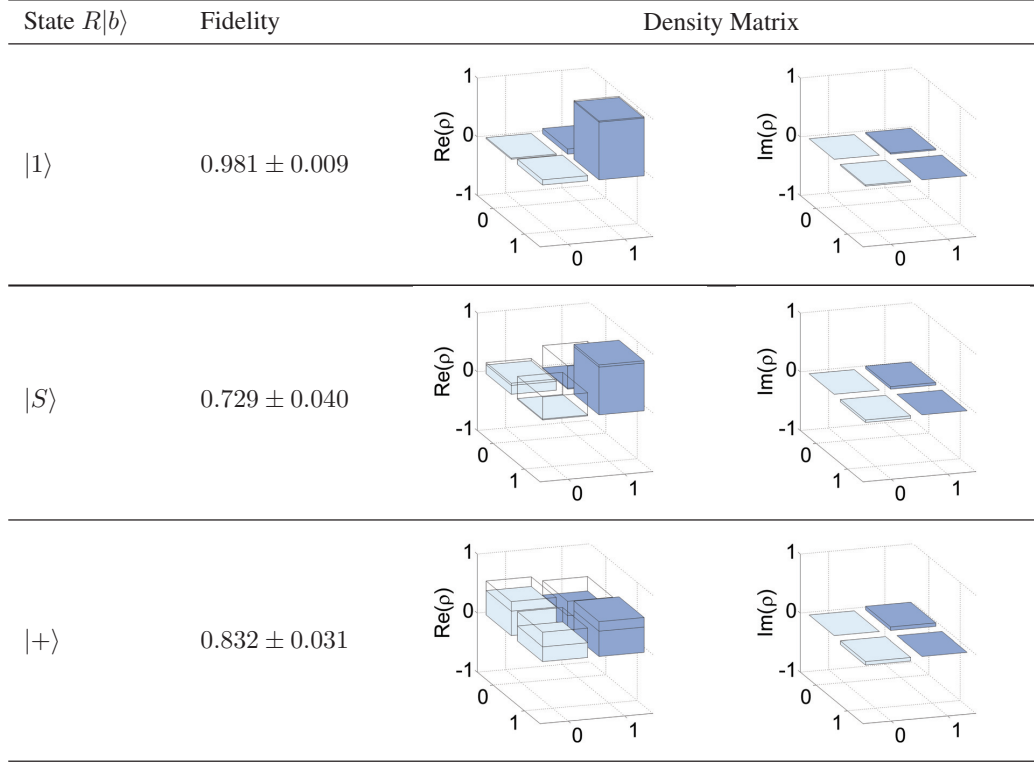


FIG. 8: The figure shows the different states $R|b\rangle$ which we have chosen as input states to our circuit.